arXiv:1612.04647v1 [cs.CV] 14 Dec 2016

UnrealStereo: A Synthetic Dataset for Analyzing Stereo Vision

Yi Zhang[1*], Weichao Qiu[2*], Qi Chen[3], Xiaolin Hu[1], and Alan Yuille[2]
[1]Tsinghua University, Beijing, China
[2]Johns Hopkins University, Baltimore, MD, USA
[3]Shanghai Jiao Tong University, Shanghai, China
{edwardz.amg, qiuwch, qchenao}@gmail.com, xlhu@mail.tsinghua.edu.cn, alan.l.yuille@gmail.com

Abstract

Stereo algorithm is important for robotics applications, such as quadcopter and autonomous driving. It needs to be robust enough to handle images of challenging conditions, such as raining or strong lighting. Textureless and specular regions of these images make feature matching difficult and smoothness assumption invalid. It is important to understand whether an algorithm is robust to these hazardous regions. Many stereo benchmarks have been developed to evaluate the performance and track progress. But it is not easy to quantize the effect of these hazardous regions. In this paper, we develop a synthetic image generation tool and build a benchmark with synthetic images. First, we manually tweak hazardous factors in a virtual world, such as making objects more specular or transparent, to simulate corner cases to test the robustness of stereo algorithms. Second, we use ground truth information, such as object mask, material property, to automatically identify hazardous regions and evaluate the accuracy of these regions. Our tool is based on a popular game engine Unreal Engine 4 and will be open-source. Many publicly available realistic game contents can be used by our tool which can provide an enormous resource for algorithm development and evaluation.

1. Introduction

Stereo algorithms benefit enormously from benchmarks [1, 17]. They provide quantitative evaluations to encourage competition and track progress. Despite the great progress over the past years, many challenges still remain unsolved, such as transparency, specularity and weakly textured regions. Image regions containing these challenges are very likely to cause the failure of an algorithm and are hazardous. These hazardous regions are small and uncommon, but critical in the real world. For example, the street light is a thin object and covers a small portion in an image, but missing it will be a disaster for autonomous driving. Raining weather is not common, but on a rainy day, significant portions of the road will be very specular. These hazardous cases are summarized by CV-HAZOP [21]. Understanding the performance of these cases is critical for avoiding real-world failures, such as the accidental crash of an autonomous car. Stereo datasets include these challenges, but these hazardous regions are usually a small portion of an image and not very common in a dataset, so they won't affect the overall number a lot. This means achieving a good performance is not necessary corresponding to solving these challenges. To better understand the robustness of stereo algorithms, we built a benchmark UnrealStereo. We use this benchmark to study the impact of four hazardous factors on state-of-the-art algorithms. These factors are specularity, no texture, disparity jumps and transparency. The severity of hazardous factors is precisely controlled by making the material of objects more specular or transparent. Hazardous regions of an image are also provided to understand their impact on performance, see Fig. 1.

We want to understand the robustness of an algorithm when it is facing extreme hazardous conditions. To achieve this, we need to provide corner cases to throughly test it. *Corner cases* are situations when parameters are at extreme levels. They are widely used to test the correctness and robustness of an algorithm. The idea of using corner cases in computer vision is attractive. But two challenges make it hard for practice:

First, it is difficult to capture corner cases. Strong lighting and rainy weather are rare and we can not control the weather to produce such images. Some factors, such as camera viewpoint and lighting, can be controlled with expensive lab setup [10, 2], but other properties such as the transparency and specularity of material are hard to control. Also, these lab setups are oversimplified and expensive. In SYNTHIA [16], synthetic images under different weather conditions are produced. Synthetic images and ground truth are available, but the virtual scenes for ren-

*Indicates equal contributions.

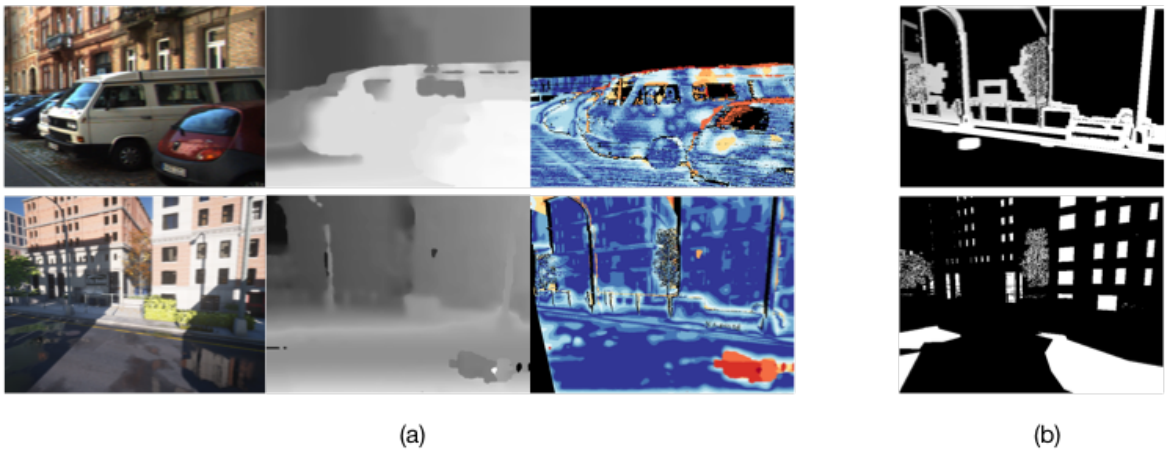

Figure 1. (a) From left to right are image, disparity estimation of MC-CNN[20] and error. The red color means large error. The first row is from KITTI[13] and the second row is our synthetic image. The error of a stereo algorithm is mainly due to hazardous factors, such as transparency and specularity. (b) Our tool can control the severity of these hazardous factors and automatically identify hazardous regions. The top row is a mask computed from object boundary, the bottom is the transparent and specular regions caused by window and wet ground. These two masks correspond to the image in the second row of (a).

dering are not publicly available. So it is not possible to do further tweaking with the rendering parameters to produce more challenging images. UnrealStereo provides tools to convert publicly available game contents into a dataset. Researchers can use images we generated or generate new corner cases to thoroughly test their algorithms.

Second, the hazardous regions of an image are unknown. When an image causes failure, we want to know which hazardous factors cause the failure. So it is useful to spot the hazardous regions of an image. But this requires information beyond disparity map. In CV-HAZOP [21], experts manually labeled the hazardous factors for demonstration, but this approach is hard to scale up. These regions can be automatically computed from information such as material property or object boundary. This information can be found in some datasets, but are absent from most stereo datasets. Lacking this information also limits the possibility of training algorithms that can utilize semantic information to solve these challenges. Computer graphics is a powerful tool for producing images and ground truth. Ground truth can be computed instead of being manually labeled or captured by special sensors. Our benchmark utilized the rich information from synthetic images to automatically identify hazardous regions of an image, which is hard for real images.

In this paper, we develop a data generation tool called UnrealStereo and use it to generate synthetic images for analyzing stereo algorithms. The contributions of UnrealStereo are: 1. Provide precise control over the hazardous factors of a scene, such as specularity and transparency of a material. 2. Provide rich ground truth such as object instance mask and material property to automatically identify hazardous regions of an image. 3. Enable the usage of rich game resources on the market to do virtual experiments. Using UnrealStereo, we first manually designed some cases which contain hazardous factors and studied the impact of these factors on algorithm performance. Then using the rich ground truth provided by UnrealStereo, we automatically spot hazardous regions of an image, such as object boundary, specular areas and analyze the performance in these regions. Using our benchmark, we can produce a report for each stereo algorithm, showing its robustness to hazardous factors. Reports for algorithms we evaluated will be included in the supplementary material.

2. Related Work

Many stereo datasets have been created for training and evaluating stereo algorithms. Middlebury stereo dataset [17] is a widely used indoor scene dataset, which provides high-resolution stereo pairs with dense disparity ground truth. KITTI stereo dataset [6, 13] is a benchmark consisting of urban video sequences while semi-dense disparity ground truth along with semantic labels are available. Due to demand of complex equipment and expensive human labor, these two real-world datasets have relatively small sizes. The larger KITTI dataset has about 400 labeled stereo pairs in total for public use. Another disadvantage of real-world dataset is the limited precision of 3D sensors and LIDAR prohibits high-quality ground truth. Recently, a large synthetic dataset [12] is rendered to leverage training deep convolutional neural networks (DCNN) end-to-end for stereo dense prediction.

Synthetic data attracts a lot of attention recently. The progress of computer graphics makes synthesizing realistic images much easier. The ability to get a large amount of images and ground truth is attractive. Synthetic data have been

used in optical flow [3], semantic segmentation [15, 16, 5], stereo [12]. Images and ground truth are provided in these datasets, but the virtual scenes are not available for various reasons. So it is not possible to render new images or change the properties of these scenes. Instead of constructing virtual scenes from scratch, we use game projects that are publicly available in the marketplace. Our tool enables generating images and ground truth from these game projects. Anyone can use our tool to tweak these virtual scenes and produce more data. Many virtual scenes constructed by visual artists in the marketplace can be used. Different from rendering images from a commercial game binary [5], the ability to access 3D models enables us to modify the scene, generate more ground truth and do various virtual experiments. Different from Sintel [3] and Flyingthing3D [12], our approach utilizes more realistic 3D models and designed for real-time rendering. Previously, OVVV [18] attempted to use synthetic data for evaluating algorithms, but there is no easy-to-use platform available.

Understanding the impact of hazardous factors is important for computer vision. CV-HAZOP [21] proposes the idea of analyzing hazardous factors in an image. Their method requires manually annotating risk factors, such as specularity area, from images, which is difficult to perform and hard to scale up. Our synthetic pipeline can automatically identify these hazardous regions, enables large-scale analysis. The ability to control the severity of hazardous factors also helps us better understand the weakness of an algorithm.

3. Method

In this section, we describe how to construct virtual worlds for synthetic image generation. We tweak the severity of harzadous factors to produce corner cases to throughly test four state of the art algorithms. Hazardous regions are computed to analyze the effect of each hazardous factor.

3.1. Construct a Virtual World

A virtual world enables a mobile agent to navigate, take actions and interact. We constructed virtual worlds to analyze the robustness of stereo algorithms to hazardous factors. Our virtual worlds are created using the popular game engine Unreal Engine 4 (UE4). Instead of building them from scratch, we bought high-quality game projects and modified them to fit our requirements. The UE4 plugin UnrealCV[14] is used to capture images and ground truth from games. It is capable of generating images, depth and object instance mask. We built our tool on top of UnrealCV. Our tool can produce accurate depth ground truth for transparent objects and support the extraction of material properties, such as the transparent or specular regions. This information enables us to provide the hazardous regions of an image in the benchmark. We also added the support for multiple cameras, so that two images can be generated simultaneously for stereo. All the game projects are available on the internet, so others can analyze other hazardous factors using our code and these projects.

To generate our test images, we use the scene editor of Unreal Engine to change the material properties of objects, making them more specular or transparent. We also put in occluder to increase the difficulty. Images are captured by moving the virtual camera according to a trajectory. The trajectory is created by two methods. First, we navigate our player and taking photos where we find interesting. Second, we use the cinematic tool of UE4 which is designed for movie production. Creating a few keyframes and intermediate frames will be interpolated. A stereo dual-camera system in standard form is put into our virtual world to produce rectified stereo pairs. Images and ground truth are generated for each frame in the camera trajectories, which is shown in Fig. 2.

Figure 2. From left to right are rendered images, object instance mask, material information (green shows transparent and red shows specular region).

Given a rectified image pair, the goal of stereo matching is to compute the disparity d for each pixel in the reference image. The disparity is defined as the difference in horizontal location of a point in the left image and its corresponding one in the right. Then the conversion between depth z and disparity d is shown in the following relation $z = \frac{fB}{d}$. where f is the focal length of the camera and B is the baseline that is the distance between the camera centers. The correctness of disparity is verified by warping the reference image according to its disparity map and comparing it with the target image.

3.2. Design Hazardous Cases for Evaluation

The task of stereo matching has several well-known difficulties including occlusion, specular surfaces, textureless regions, semi-transparent objects and regions with many jumps in disparity, i.e. frequent disparity discontinuities. We call these factors hazard as mentioned in [21]. Local stereo methods [7, 11] suffer from matching ambiguities in weakly textured, specular and semi-transparent regions. Global methods [9, 19, 20, 4], which impose smoothness constrains on neighboring pixels or superpixels, have the ability to overcome some of the above problems. Special efforts have been made to solve these difficulties. The authors of [8] leverage semantic informantions and 3D CAD

models to resolve stereo ambiguities brought by specularity or no texture. An end-to-end trained DCNN based algorithm [12] performs well on specular regions of KITTI stereo 2015 [13] after finetuning on the training set. They exploit a fully convolutional siamese architecture to predict dense disparity map at a single forward pass.

(a) Specularity (b) No texture

(c) Disparity jumps (d) Transparency

Figure 3. From (a) to (d) are cases we designed to test algorithms, includes specularity, no texture, disparity jump and transparency

To conduct an evaluation of these algorithms on the mentioned challenges, we manually designed and evaluate popular stereo algorithms on four hazardous factors i.e. specularity, transparency, no texture and rapid disparity jumps. While more than one hazardous factor above often cooccur in real world situations, our method enables to investigate the effect of each factor separately. All the four cases are established in a virtual living room while different objects are present. Fig. 3 shows the outlines of the four factors. These cases are important due to the following reasons. Firstly, these difficult cases are very common in the real world and often encountered by either indoor or outdoor stereo vision applications. Results on available stereo datasets such as KITTI [6, 13] and Middlebury [17] show a great deteriorated performance for all the algorithms on these cases. Secondly, it is convenient for researchers to analyze the performance of their methods on simple cases where only one factor is present. Furthermore, choosing the solutions for related applications entails a comparison of stereo algorithms when dealing with these hazardous cases. Specifically, in some applications where specular or semi-transparent objects appear frequently, the method with the lower error on such regions would be more likely to be chosen.

In the specularity challenge(Fig. 3(a)), we obtain the specular effect by diminish the roughness of specific materials. The major specular object is the screen of a TV towards which test stereo pairs are obtained from various viewpoints and distances. In the no texture challenge(Fig. 3(b)), the wall and the ceiling in the room are made textureless because they are the most common textureless objects in real world. To achieve texturelessness while keep reality, we do not directly remove the material of the walls but increase the smoothness of their material. A set of ten test image pairs are collected covering different orientations to the surrounding textureless walls. In the disparity jumping case(Fig. 3(c)), we place in the scene objects such as bamboos, fences and plants of various sizes and poses, which easily form many disparity discontinuities distributed within a small region. In the transparency case(Fig. 3(d)), we placed a transparent sliding wall in a room.

Not only can we construct ideal cases to evaluate stereo algorithm under single difficult condition, for specularity, semi-transparency and no texture, it is also possible to investigate how their performance degrades as the extent of challenge increases. Unlike in real world, adjusting the attribute of materials in virtual world is convenient and costless. This kind of analysis enables us to compare the robustness of stereo matching algorithms, thus making better choice on method selection depending on specific application. Fig. 4 shows the result of making the TV monitor more specular. It is shown that the state-of-art method MC-CNN [20] failed to understand the concept of TV and made a wrong estimation. Consequently, if a quadcopter mistakenly considers this TV as a window and tries to fly through it, this will be a failure. Details and results for evaluation on these cases are presented in Section 4.1.

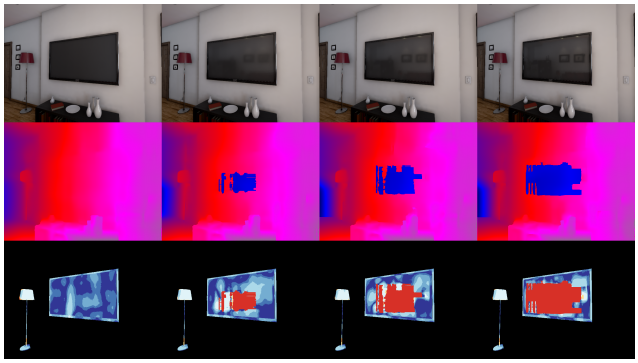

Figure 4. Different levels of specularity, from top to bottom are input image, disparity estimation and error compared with ground truth, the error is only computed for the TV and the floor lamp.

3.3. Automatic Hazardous Region Discovery

It is important to pay more attention to hazardous regions of an image, because these regions are most likely to cause the failure of an algorithm. Manually designed hazardous cases are important for understanding an algorithm, but the size of manually designed data is not large enough.

On the other hand, due to the popularity of virtual reality, there are a lot of high quality virtual environments. These virtual environments can be purchased with a fair price (less than $100) or even free. These virtual environments can produce large amount of images. Moreover, we can also access other information, such as object instance mask, material property. We use these information to compute hazardous regions and to evaluate stereo algorithms on these hazardous regions.

Our rendering process produces extra information beyond depth information. These information includes: object instance mask, specular region, transparent region. Using these extra information, we can locate these hazardous regions automatically. Specifically, we add new focus to stereo method evaluation by providing binary masks for the these regions: 1. object boundary, 2. weakly textured material, 3. specular material, 4. transparent material. These regions are related to what we mention in Section 3.2 respectively. To evaluate disparity jumps hazard, boundary regions are exploited because theoretically frequent disparity discontinuities add difficulty to the computation of initial estimation, thus methods within traditional framework [19, 4, 20] fail on these regions especially at the boundary. Fig. 5 shows an example of these masks. All of above masks are generated automatically from the object segmentation ground truth. The object boundary mask is obtained by first computing a binary edge map from the object segmentation and dilating it with a radius of r_d to include the vicinity of edges. For each object, we annotate its material information only once, before rendering process, then no more human effort is required to obtain corresponding masks.

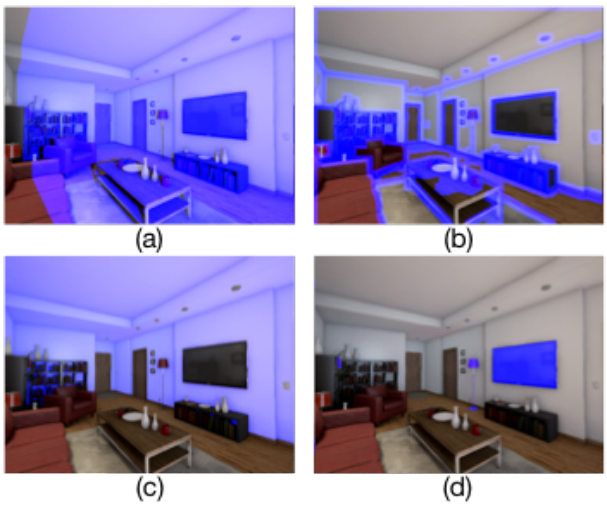

Figure 5. Mask that we compute from object mask and material property. From (a) to (d) are: mask for non-occluded region, boundary region, textureless region and specular region.

4. Experiment

We use three publicly available game scenes in our experiment. Screenshot of these games can be seen in Fig. 6. The comparison with other stereo datasets is shown in Table. 1. Besides depth and disparity ground truth, we provide extra information, such as: object segmentation mask and material properties. Another unique feature of our dataset is the properties of our virtual worlds can be tweaked with our tool and more challenging images can be produced. Instead of just providing an image dataset with fixed number of images, we provide a synthetic image generation tool. This tool can be used to design new hazardous cases, generate more images. More game scenes from the marketplace can be used in experiment.

	Density	Scene Type	Extra Info.	Render on Demand
Middlebury[17]	dense	laboratory		
KITTI2012[6]	50%	outdoor	seg.	
KITTI2015[13]	50%	outdoor	seg.	
Virtual KITTI[5]	dense	outdoor	seg.	
FlyingThings3D[12]	dense	toy, outdoor	seg., material	
Ours	dense	indoor, outdoor	seg., material	√

Table 1. Comparison of stereo benchmarks and datasets. Our dataset and data generation tool can provide more information, such as: object segmentation mask, material properties. It can also produce new challenging images from a virtual world

We test local method CoR [4], global methods on pixel-level MC-CNN [20] and superpixel-level SPS-St [19] as well as DCNN based methods DispNet [12]. Implementation from the authors of these methods are adopted. For completeness, we use two versions of [20], one is fast, the other is accurate. We also include a variant of [12] DispNetC which performs better on KITTI benchmark [6]. For the weights of the MC-CNN [20], we adopt the models used in their submission to KITTI. While for the end-to-end model [12], weights trained on the authors' original synthetic dataset are used. We also tried the fine-tuned model, but the performance is lower. In our experiments, we choose two commonly adopted error metrics, i.e., the end-point error which takes average of the absolute difference from ground truth and the 3 pixel error which is the percentage of pixels which deviate from ground truth more than 3 pixel. The 3 pixel error metric is proposed to cover calibration and laser measurement errors in real-world datasets and we use it here to compare results with them. The following evaluations are mainly focused on end-point error.

4.1. Hazardous Cases Evaluation

We choose 10 different viewpoint for each of the hazardous cases we designed, i.e. specular, semi-tranparent, textureless, and disparity jumps, covering both fronto-parallel and slanted surfaces. At each viewpoint of hazardous scenes except disparity jumps case, we start from

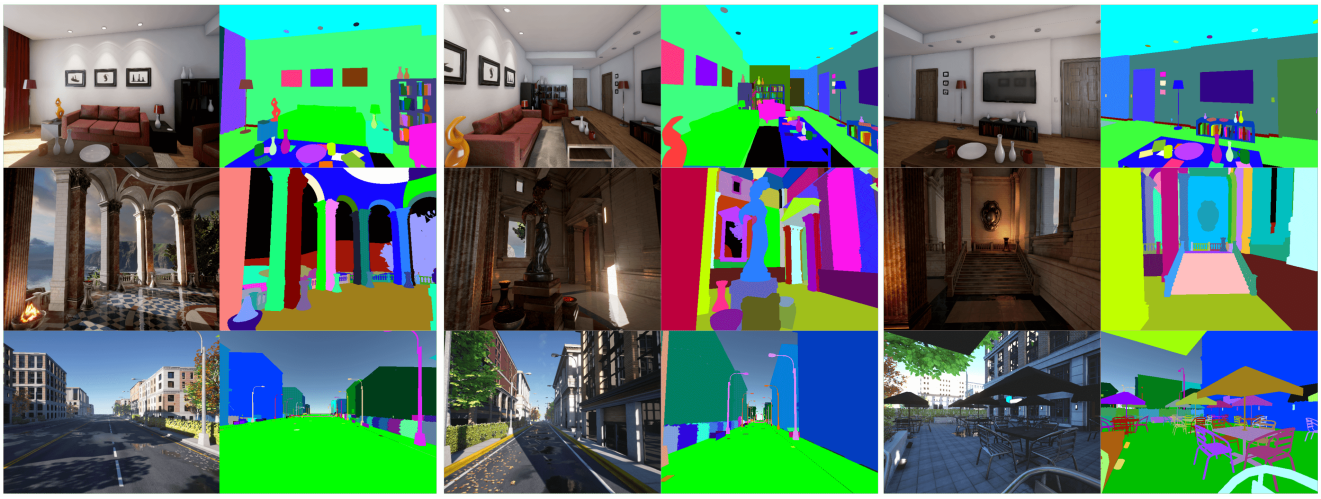

Figure 6. The virtual scenes we use in our experiments, from top to bottom are RealisticRendering, SunTemple, UrbanCity. These virtual scenes can be purchased of freely downloaded from Unreal Engine marketplace.

the easiest parameter settings that are roughest, opaque or well-textured and adjust the corresponding parameter step by step to increase the extent of hazard, resulting in 8 different levels of corresponding hazard per viewpoint and a total od 80 stereo image pairs. In specularity, semi-transparent and textureless cases, only regions labeled as these types are evaluated. For all the cases, we evaluate non-occluded regions because occlusion will introduce new problems to the task while we focus on analyzing one factor at a time. Table. 2 shows the performance on extreme condition. Fig. 7 displays the comparison of degradation of state-of-the-art stereo algorithms.

Type of Hazard: As shown in Table 2, the slanted plane model based method SPS-St [19] performs best in specular case, closely followed by CoR [4] which is not purely local and reasons with overlapping multi-scale set of regions. As for high transparency, tremendously high test errors are shown, which means the little cue from the blur and specularity on transparent surfaces is insufficient for these method to spot the existence of a front surface. The success of DCNN models DispNet and DispNetC [12] on no texture and disparity jumping cases is evident. It is because they have large receptive fields to incorporate context information and do not explicitly impose smoothness constraints. The state-of-the-art MC-CNN [20] has good overall performance, but is less robust in extremely hazardous scenes - it easily suffers from ambiguity brought by no texture, specularity and transparency.

Levels of Difficulty: A qualitative result of specularity is shown in Fig. 4. Both the qualitative and quantitative results in Fig. 7 demonstrate a degradation of stereo algorithms as the level of difficulty increases. In no texture case, end-to-end DCNN based algorithms [12] have higher endurance for textureless objects, which is conceivably attributed to the ability of DCNN to utilize larger context information. MC-CNN [20] produces accurate result when texture is not so weak. However, as the texture become increasingly weak, when the matching network fails to produce a unambiguous data term, its local smoothness term, which encourages fronto-parallel surfaces, would leads to huge error in these regions. For specular regions, both SPS-St [19] and CoR [4] show higher robustness over other methods, indicating that the large support regions they use to some extent reduces their vulnerability to matching ambiguities. Results on transparency hazard are similar to those on specularity, because they bring the same problem for matching.

4.2. Automatic Hazardous Region Discovery

We utilize the game scenes to generate four binocular video sequences with disparity ground truth and hazardous regions, each of which has approximately 500 frames and includes several kinds of hazardous cases. These sequences are described as follows:

RealisticRendering: Indoor environment of a living room with common furniture, plants as well as specular surface of a TV screen.

UnrealMiddlebury: Indoor scene similar to Middlebury [17]. Objects of various shapes are placed together, which increases the difficulty with more occlusion and disparity discontinuity.

SunTemple: Large indoor environment with few objects and no particular hazard being present.

UrbanCity: A KITTI like urban scene sequence. The water on the roads simulates the driving condition after a heavy rain.

Stereo algorithms are tested on each video sequence and errors averaged within each sequence are reported. Table 3 shows the evaluation based on harzardous regions. The re-

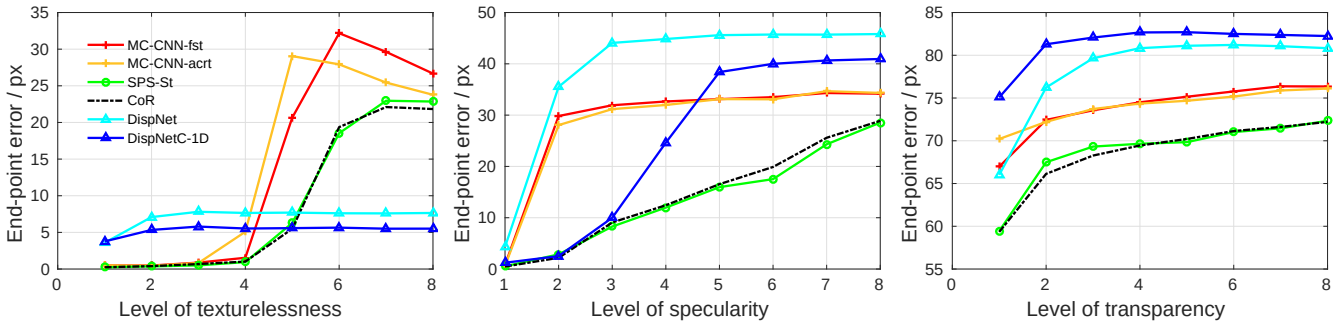

Figure 7. The influence of texturelessness, specularity and transparency at different levels. The performance of stereo algorithms is evaluated in terms of end-point error. Larger number of level represents is controlled by parameters for corresponding materials. Each data point represents an average over 10 different viewpoint.

	Specular		Textureless		Disparity Jumps		Tranparency	
	EPE	>3px	EPE	>3px	EPE	>3px	EPE	>3px
CoR[4]	26.4	56.5%	8.35	37.5%	7.26	23.11%	**59.34**	76.93%
SPS-St[19]	**25.2**	**53.7%**	8.98	37.1%	7.50	23.89%	59.44	**76.13%**
MC-CNN-fst[20]	30.3	65.7%	7.97	**24.1%**	7.37	22.18%	66.97	86.10%
MC-CNN-acrt[20]	30.8	66.9%	7.97	25.8%	7.00	**21.51%**	70.21	86.57%
DispNet[12]	36.7	92.3%	**6.59**	45.0%	6.72	38.92%	66.01	85.74%
DispNetC[12]	36.3	88.4%	8.13	55.4%	**6.18**	34.71%	75.08	94.28%

Table 2. Performance on extreme hazardous cases. Result in end-point error (EPE) and 3 pixel error are presented for specular, no texture, disparity jumping and transparent hazard. Except for disparity jumps, only regions annotated as corresponding hazard are evaluated.

sults indicate that the boundary and specular regions obtained automatically are indeed hazardous to stereo algorithms.

For overall performance in full and non-occluded regions, results on UrbanCity is consistent with the ranking on KITTI [6] both in terms of end-point and 3 pixel error metric, which indicates the authenticity of our data rendered from virtual scenes. Moreover, for specular regions, results on RealisticRendering and UrbanCity also agree with those on KITTI, except for the performance of DispNetC [12]. For Middlebury like scene in UnrealMiddlebury, significant deterioration of performance is observed. Therefore, the frequent disparity jumps and the transformation of shape due to viewpoint change significantly influences the difficulty of stereo matching. In large indoor environment as SunTemple, the lack of hazards results in relatively low errors. The top ranking method on KITTI MC-CNN [20] outperforms other method as expected.

On the boundary, errors higher than non-occluded regions are observed for all algorithms, which indicates that boundary region is a hazardous factor. This suggest that current stereo methods have difficulty predicting edge-preserved disparity maps, probably due to unsuccessful matching or the failure of segmentation to align well with the boundaries.

Although generated in much more complex environment, most of the results on hazardous regions still coincide with the manually designed cases, i.e. robustness of stereo algorithms varies under different levels of hazardous factors. For specularity, SPS-St [19] and CoR [4] outperform other methods on RealisticRendering and UrbanCity, while MC-CNN [20] performs best on SunTemple. From Fig. 7 we find that they coincide respectively with high and low specularity conditions. For textureless materials, results on UrbanCity match the weakest textured case, where end-to-end DCNN based models perform best, while the other three could be interpreted as cases with evident texture.

SPS-St [19] and CoR [4] have similar performance both in designed case and rendered natural scenes. The latter uses the former to compute an initial semi-dense estimation of disparity and a multi-scale aggregation mechanism to further refine the result. In scenes with rich hazardous factors, e.g. RealisticRendering, UnrealMiddlebury and UrbanCity, adopting this mechanism results in better accuracy, while in simple scenes, e.g. SunTemple, no such gain in performance is observed.

The performance of DCNN based models [12] drops significantly on these rendered video sequences except for SunTemple. Since DCNN models highly hinges on training data, whether training data includes hazardous cases could determine the respective performance. Therefore, for DCNN models, finetuning becomes important when transfer from one dataset to another.

The state-of-the-art pixel-level global method MC-

	Full		Non-Occluded		Boundary		Textureless		Specular	
	EPE	>3px	EPE	>3px	EPE	>3px	EPE	>3px	EPE	>3px
RealisticRendering										
CoR[4]	**2.23**	**10.52%**	**1.05**	5.70%	**1.25**	7.10%	**0.86**	5.08%	**4.59**	24.18%
SPS-St[19]	2.56	11.68%	1.12	5.70%	1.27	6.91%	0.92	5.32%	4.90	**21.10%**
MC-CNN-fst[20]	3.11	12.24%	1.42	5.65%	1.57	7.09%	1.18	4.72%	7.22	29.43%
MC-CNN-acrt[20]	3.01	11.79%	1.31	**5.35%**	1.32	**6.07%**	1.01	**4.55%**	7.19	29.56%
DispNet[12]	4.71	32.40%	3.58	28.12%	2.97	23.87%	4.00	33.14%	9.02	51.13%
DispNetC[12]	4.71	28.66%	3.58	24.62%	2.58	16.67%	4.58	33.66%	7.91	46.10%
UnrealMiddlebury										
CoR[4]	13.34	**24.69%**	6.86	12.78%	8.96	20.23%	1.34	3.71%	34.80	62.01%
SPS-St[19]	14.35	25.08%	6.85	12.88%	8.91	20.14%	1.49	4.29%	29.02	57.42%
MC-CNN-fst[20]	13.32	26.11%	**6.29**	12.44%	**7.94**	19.02%	1.43	3.75%	22.91	45.59%
MC-CNN-acrt[20]	13.62	25.54%	6.48	**11.88%**	7.97	**18.03%**	**1.30**	**2.78%**	**17.07**	**37.24%**
DispNet[12]	13.38	48.08%	7.60	38.81%	10.95	47.96%	3.88	42.30%	26.27	88.58%
DispNetC[12]	**12.24**	41.89%	6.77	32.07%	9.26	38.95%	4.69	50.81%	21.70	81.78%
SunTemple										
CoR[4]	1.18	6.40%	0.64	3.23%	0.93	5.56%	0.99	6.23%	0.65	3.46%
SPS-St[19]	1.13	6.04%	0.54	2.63%	0.82	4.66%	0.88	5.36%	0.71	3.89%
MC-CNN-fst[20]	**1.01**	5.30%	**0.47**	1.85%	0.81	4.18%	**0.83**	**4.67%**	0.64	2.86%
MC-CNN-acrt[20]	**1.01**	**4.88%**	0.49	**1.54%**	**0.80**	**3.44%**	0.91	4.76%	**0.60**	**2.56%**
DispNet[12]	1.81	13.96%	1.45	10.98%	1.63	12.61%	1.98	15.67%	2.29	23.82%
DispNetC[12]	1.44	9.56%	1.08	6.72%	1.33	8.87%	1.57	12.18%	1.62	13.63%
UrbanCity										
CoR[4]	3.25	10.72%	2.20	6.79%	2.40	10.10%	11.38	30.88%	**1.79**	**5.24%**
SPS-St[19]	3.64	11.22%	2.49	7.06%	2.62	9.90%	12.65	29.25%	2.10	6.78%
MC-CNN-fst[20]	**3.05**	**10.61%**	**1.96**	**6.46%**	**2.02**	8.37%	8.93	26.77%	2.33	7.23%
MC-CNN-acrt[20]	3.22	10.87%	2.12	6.66%	2.18	**8.28%**	9.15	**25.13%**	2.58	8.08%
DispNet[12]	3.86	24.10%	3.15	20.18%	3.24	22.64%	7.70	44.58%	3.57	29.76%
DispNetC[12]	3.44	18.80%	2.66	14.62%	2.71	16.60%	**7.32**	36.52%	3.04	20.58%

Table 3. Performance of state-of-the-art stereo algorithms on rendered sequences. This table shows the disparity errors averaged over all test images in each video sequence. Errors in total image (Full), non-occluded regions, boundaries, textureless materials and specular regions in terms of end-point error (EPE) and ≥ 3 pixel error are evaluated by applying the masks proposed in Section 3.3.

CNN [20] receives best overall performance. However, it easily suffer from ambiguities brought by specularity and transparency. The convolutional neural network (CNN) for patch matching provide better initial disparity estimation on unambiguous regions, meanwhile, make the whole algorithm more vulnerable to specularity and transparency.

The relation of the two error metrics is worth noticing. End-point error (EPE) measures the error in average, while 3 pixel error measures the percentage of incorrect pixels. They are not in agreement in some cases, which could reveal characteristics of some algorithms. Specifically, for DCNN based method, significantly higher 3 pixel errors is observed. The L1 loss function used in the training process of DCNN models encourages lower EPE, which could account for the discrepancy.

5. Conclusion

In this paper, we present a data generation tool UnrealStereo to generate synthetic images to create a stereo benchmark. We analyzed the effect of four hazardous factors on state-of-the-art algorithms. They are specularity, no texture, disparity jumps and transparency. Each factor is precisely controlled to see its impact. They are tuned to an extreme level to produce corner cases. We also tested these algorithms on three realistic virtual scenes. Hazardous regions of each image are automatically labeled from ground truth such as object mask and material properties. We found that the state-of-the-art method MC-CNN [20] outperforms others in general, but lack robustness in hazardous cases especially specular regions. DCNN based method [12] exhibits interesting properties due to the awareness of larger context.

Our data generation tool can be used to produce more challenging images and compatible with publicly available high-quality game models. This makes our tool capable for many applications. In our future work, we will extend our platform to include more hazardous factors such as the ratio of occlusion and analyze more computer vision problems. It is also interesting to explore the rich ground truth we gen-

erate, such as object mask and material properties. This semantic information enables the development of stereo algorithms that utilizes high-level knowledge.

References

[1] S. Baker, D. Scharstein, J. Lewis, S. Roth, M. J. Black, and R. Szeliski. A database and evaluation methodology for optical flow. *International Journal of Computer Vision*, 92(1):1–31, 2011. 1

[2] A. Borji, S. Izadi, and L. Itti. ilab-20m: A large-scale controlled object dataset to investigate deep learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2221–2230, 2016. 1

[3] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In *European Conference on Computer Vision*, pages 611–625. Springer, 2012. 3

[4] A. Chakrabarti, Y. Xiong, S. J. Gortler, and T. Zickler. Low-level vision by consensus in a spatial hierarchy of regions. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4009–4017. IEEE, 2015. 3, 5, 6, 7, 8

[5] A. Gaidon, Q. Wang, Y. Cabon, and E. Vig. Virtual worlds as proxy for multi-object tracking analysis. *arXiv preprint arXiv:1605.06457*, 2016. 3, 5

[6] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 3354–3361. IEEE, 2012. 2, 4, 5, 7

[7] A. Geiger, M. Roser, and R. Urtasun. Efficient large-scale stereo matching. In *Asian conference on computer vision*, pages 25–38. Springer, 2010. 3

[8] F. Guney and A. Geiger. Displets: Resolving stereo ambiguities using object knowledge. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4165–4175, 2015. 3

[9] H. Hirschmuller. Accurate and efficient stereo processing by semi-global matching and mutual information. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 2, pages 807–814. IEEE, 2005. 3

[10] Y. LeCun, F. J. Huang, and L. Bottou. Learning methods for generic object recognition with invariance to pose and lighting. In *Computer Vision and Pattern Recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE Computer Society Conference on*, volume 2, pages II–97. IEEE, 2004. 1

[11] Z. Ma, K. He, Y. Wei, J. Sun, and E. Wu. Constant time weighted median filtering for stereo matching and beyond. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 49–56, 2013. 3

[12] N. Mayer, E. Ilg, P. Häusser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 3, 4, 5, 6, 7, 8

[13] M. Menze and A. Geiger. Object scene flow for autonomous vehicles. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3061–3070, 2015. 2, 4, 5

[14] W. Qiu and A. Yuille. Unrealcv: Connecting computer vision to unreal engine. *arXiv preprint arXiv:1609.01326*, 2016. 3

[15] S. R. Richter, V. Vineet, S. Roth, and V. Koltun. Playing for data: Ground truth from computer games. In *European Conference on Computer Vision*, pages 102–118. Springer, 2016. 3

[16] G. Ros, L. Sellart, J. Materzynska, D. Vazquez, and A. M. Lopez. The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3234–3243, 2016. 1, 3

[17] D. Scharstein and R. Szeliski. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International journal of computer vision*, 47(1-3):7–42, 2002. 1, 2, 4, 5, 6

[18] G. R. Taylor, A. J. Chosak, and P. C. Brewer. Ovvv: Using virtual worlds to design and evaluate surveillance systems. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2007. 3

[19] K. Yamaguchi, D. McAllester, and R. Urtasun. Efficient joint segmentation, occlusion labeling, stereo and flow estimation. In *European Conference on Computer Vision*, pages 756–771. Springer, 2014. 3, 5, 6, 7, 8

[20] J. Zbontar and Y. LeCun. Computing the stereo matching cost with a convolutional neural network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1592–1599, 2015. 2, 3, 4, 5, 6, 7, 8

[21] O. Zendel, M. Murschitz, M. Humenberger, and W. Herzner. Cv-hazop: Introducing test data validation for computer vision. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2066–2074, 2015. 1, 2, 3